arXiv:1610.01687v1 [cs.GT] 5 Oct 2016

Sampled Fictitious Play is Hannan Consistent

Zifan Li
zifanli@umich.edu

Ambuj Tewari
tewaria@umich.edu

October 7, 2016

Abstract

Fictitious play is a simple and widely studied adaptive heuristic for playing repeated games. It is well known that fictitious play fails to be Hannan consistent. Several variants of fictitious play including regret matching, generalized regret matching and smooth fictitious play, are known to be Hannan consistent. In this note, we consider sampled fictitious play: at each round, the player samples past times and plays the best response to previous moves of other players at the sampled time points. We show that sampled fictitious play, using Bernoulli sampling, is Hannan consistent. Unlike several existing Hannan consistency proofs that rely on concentration of measure results, ours instead uses anti-concentration results from Littlewood-Offord theory.

Keywords: adaptive heuristics, learning, repeated games, Hannan consistency, fictitious play

1 Introduction

In the setting of repeated games played in discrete time, the (unconditional) regret of a player, at any time point, is the difference between the payoffs she would have received had she played the best, in hindsight, constant strategy throughout, and the payoffs she did in fact receive. Hannan [1957] showed the existence of procedures with a "no-regret" property: procedures for which the average regret per time goes to zero for a large number of time points. His procedure was a simple modification of fictitious play: random perturbations are added to the cumulative payoffs of every strategy so far and the player picks the strategy with the largest perturbed cumulative payoff. No regret procedures are also called "universally consistent" [Fudenberg and Levine, 1998, Section 4.7] or "Hannan consistent" [Cesa-Bianchi and Lugosi, 2006, Section 4.2].

It is well known that adding perturbations to the cumulative payoffs before computing best response to other players' previous moves is crucial to achieve Hannan consistency. Without perturbations, the procedure becomes identical with fictitious play, which fails to be Hannan consistent [Cesa-Bianchi and Lugosi, 2006, Exercise 3.8]. Besides Hannan's modification, other variants of fictitious play are also known to be Hannan consistent, including (unconditional) regret

matching, generalized (unconditional) regret matching and smooth fictitious play (for an overview, see Hart and Mas-Colell [2013, Section 10.9]).

In this note, we consider another variant of fictitious play, namely sampled fictitious play. Here, the player samples past time points using some (randomized) sampling scheme and plays best response to the plays of the other players restricted to the set of sampled time points. Sampled fictitious play has been considered by other authors in the context of evolutionary games [Kaniovski and Young, 1995], the game of matching pennies [Gilliland and Jung, 2006], and games with identical payoffs [Lambert III et al., 2005]. To the best of our knowledge, it is not known whether sampled fictitious play is Hannan consistent. The purpose of this note is to show that it is indeed Hannan consistent when used with a natural sampling scheme, namely Bernoulli sampling.

2 Preliminaries

Consider a game in strategic form where M is the number of players, S_i is the set of strategies for player i, and $u_i : \prod_{j=1}^{M} S_i \to \mathbb{R}$ is the payoff function for player i. For simplicity assume that the payoff functions of all players are $[0,1]$ bounded. We also assume the number of strategies is the same for each player and that $S_i = \{1, \dots, N\}$. Let $S = \prod_{i=1}^{M} S_i$ be the set of M-tuples of player strategies. For $s = (s_i)_{i=1}^{M} \in S$, we denote the strategies of players other than i by $s_{-i} = (s_j)_{1 \le j \le M, j \ne i}$.

The game is played repeatedly over (discrete) time $t = 1, 2, \dots$. A learning procedure for player i is a procedure that maps the history $h_{t-1} = (s_\tau)_{\tau=1}^{t-1}$ of plays just prior to time t, to a strategy $s_{t,i} \in S_i$. The learning procedure is allowed to be randomized, i.e., player i has access to a stream of i.i.d. random variables $\epsilon_1, \epsilon_2, \dots$ and she is allowed to use $\epsilon_1, \dots, \epsilon_{t-1}$, in addition to h_{t-1}, to choose $s_{t,i}$. Player i's regret at time t is defined as

$$\mathcal{R}_{t,i} = \max_{k \in S_i} \sum_{\tau=1}^{t} u_i(k, s_{\tau,-i}) - \sum_{\tau=1}^{t} u_i(s_\tau).$$

This compares the player's cumulative payoff with the payoff she could have received had she selected the best constant (over time) strategy k with knowledge of the other players' moves.

A learning procedure for player i is said to be *Hannan consistent* iff

$$\limsup_{t \to \infty} \frac{\mathcal{R}_{t,i}}{t} \le 0 \qquad \text{almost surely.}$$

Hannan consistency is also known as the "no-regret" property and as "universal consistency". The term "universal" refers to the fact that the regret per time goes to zero irrespective of what the other players do.

Fictitious play is a (deterministic) learning procedure where player i plays best response to the plays of the other players so far. That is,

$$s_{t,i} \in \arg\max_{k=1}^{N} \sum_{\tau=1}^{t-1} u_i(k, s_{\tau,-i}).$$

Fictitious play is not Hannan consistent. However, consider the following modification of fictitious play, called *sampled fictitious play*. At time t, player randomly selects a subset $\mathbb{S}_t \subseteq \{1, \ldots, t-1\}$ of previous time points and plays best response to the other players' moves only over $\mathbb{S}_t$. That is,

$$s_{t,i} = \begin{cases} s_{t-1,i} & \text{if } s_{t-1,i} \in \arg\max_{k=1}^{N} \sum_{\tau \in \mathbb{S}_t} u_i(k, s_{\tau,-i}) \\ \in \arg\max_{k=1}^{N} \sum_{\tau \in \mathbb{S}_t} u_i(k, s_{\tau,-i}) & \text{otherwise} \end{cases}. \tag{1}$$

The first case above ensures that if the previous time point's strategy $s_{t-1,i}$ is still the best response at the current round, the player does not switch to a new strategy. In the second case, if multiple strategies achieve the maximum, then the tie is broken uniformly at random. Also, if $\mathbb{S}_t$ turns out to be empty (an event that happens with probability exactly $2^{-(t-1)}$ under the Bernoulli sampling described below), we adopt the convention that the argmax above includes all N strategies.

In this note, we consider *Bernoulli sampling*, i.e., any particular round $\tau \in \{1, \ldots, t-1\}$ is included in $\mathbb{S}_t$ independently with probability $1/2$. More specifically, if $\epsilon_1^{(t)}, \ldots, \epsilon_{t-1}^{(t)}$ are i.i.d. symmetric Bernoulli (or Rademacher) random variables taking values in $\{-1, +1\}$, then

$$\mathbb{S}_t = \{\tau \in \{1, \ldots, t-1\} \,:\, \epsilon_\tau^{(t)} = +1\} \tag{2}$$

and therefore,

$$\sum_{\tau \in \mathbb{S}_t} u_i(k, s_{\tau,-i}) = \sum_{\tau=1}^{t-1} \frac{(1 + \epsilon_\tau^{(t)})}{2} u_i(k, s_{\tau,-i}).$$

Note that the procedure defined by the combination of (1) and (2) is completely parameter free, i.e., there is no tuning parameter that has to be carefully tuned in order to obtain desired convergence properties.

3 Result and Discussion

Our main result is the following.

Theorem 3.1. *Sampled fictitious play* (1) *with Bernoulli sampling* (2) *is Hannan consistent.*

Before we move on to the proof, a few remarks are in order.

Rate of convergence Our proof gives the rate of convergence of (expected) average regret as $O(N^2\sqrt{\log\log t/t})$ where the constant hidden in $O(\cdot)$ notation is small and explicit. It is known that the optimal rate is $O(\sqrt{\log N/t})$ [Cesa-Bianchi and Lugosi, 2006, Section 2.10]. Therefore, our rate of convergence is almost optimal in t but severely suboptimal in N. This raises several interesting questions. What is the best bound possible for Sampled Fictitious Play with Bernoulli sampling? Is there a sampling scheme for which Sampled Fictitious Play procedure achieves the optimal rate of convergence?

Asymmetric probabilities Instead of using symmetric Bernoulli probabilities, we can choose $\epsilon_\tau^{(t)}$ such that $P(\epsilon_\tau^{(t)} = +1) = \alpha$. As $\alpha \to 1$, the learning procedure becomes fictitious play and as $\alpha \to 0$, it selects strategies uniformly at random. Therefore, it is natural to expect that the regret bound will blow up near the two extremes of $\alpha = 1$ and $\alpha = 0$. We can make this intuition precise but only for $\{0, 1\}$-valued payoffs (instead of $[0, 1]$-valued). See Appendix B in the supplementary material.

Follow the perturbed leader Note that

$$\arg\max_{k=1}^{N} \sum_{\tau=1}^{t-1} \frac{(1+\epsilon_\tau^{(t)})}{2} u_i(k, s_{\tau,-i}) = \arg\max_{k=1}^{N} \sum_{\tau=1}^{t-1} u_i(k, s_{\tau,-i}) + \sum_{\tau=1}^{t-1} \epsilon_\tau^{(t)} u_i(k, s_{\tau,-i}).$$

Therefore, we can think of sampled fictitious play as adding a random perturbation to the expression that fictitious play optimizes. Such algorithms are referred to as "follow the perturbed leader" (FPL) in the computer science literature ("fictitious play" is known as "follow the leader"). This family was popularized by Kalai and Vempala [2005]. Closer to this paper are the FPL algorithms of Devroye et al. [2013] and van Erven et al. [2014]. However, none of these papers considered sampled fictitious play.

Extension to unconditional (or internal) regret In this paper we focus on conditional (or external) regret. Other notions of regret, especially unconditional (or internal) regret can also be considered. Internal regret measures the worst regret, over $N(N-1)$ choices of $k \neq k'$, of the form "every time strategy k was picked, strategy k' should have been picked instead". There are generic conversions [Stoltz and Lugosi, 2005, Blum and Mansour, 2007] that will convert any learning procedure with small external regret to one with small internal regret. These conversion, however, require access to the probability distribution over strategies at each time point. This probability distribution can be approximated, to arbitrary accuracy, by making the choice of the strategy in (1) multiple times each time selecting the random subset $\mathbb{S}_t$ independently. However, doing so and using a generic conversion from external to internal regret will lead to a cumbersome overall algorithm. It will be nicer to design a simpler sampling based learning procedure with small internal regret.

4 Proof of the Main Result

We break the proof of our main result into several steps. The first and third steps involve fairly standard arguments in this area. Our main innovations are in step two.

4.1 Step 1: From Regret to Switching Probabilities

In this step, we assume that players other than player i (the "opponents") are *oblivious*, i.e., they do not adapt to what player i does. Mathematically, this means that the sequence $s_{t,-i}$ is fixed and does not depend on the moves $s_{t,i}$ of player i. Since player i is fixed for the rest of the proof, we will not carry the index i in our notation further. Let the vector $g_t \in [0, 1]^N$ be defined as

$g_{t,k} = u_i(k, s_{t,-i})$ for $k \in \{1, \ldots, N\}$. Moreover, we denote player i's move $s_{t,i}$ as k_t. With this notation, regret at time T equals

$$\mathcal{R}_T = \max_{k=1}^{N} \sum_{t=1}^{T} g_{t,k} - \sum_{t=1}^{T} g_{t,k_t}.$$

In this step, we will look at expected regret. Because the opponents are oblivious, this equals

$$\mathbb{E}\left[\mathcal{R}_T\right] = \max_{k=1}^{N} \sum_{t=1}^{T} g_{t,k} - \mathbb{E}\left[\sum_{t=1}^{T} g_{t,k_t}\right] = \max_{k=1}^{N} \sum_{t=1}^{T} g_{t,k} - \sum_{t=1}^{T} \mathbb{E}\left[g_{t,k_t}\right].$$

Therefore, we do not have to draw a fresh sample $\epsilon_1^{(t)}, \ldots, \epsilon_{t-1}^{(t)}$ at time t. Instead we fix a single stream $\epsilon_1, \epsilon_2, \ldots$ of i.i.d. Radamacher random variables and set $(\epsilon_1^{(t)}, \ldots, \epsilon_{t-1}^{(t)}) = (\epsilon_1, \ldots, \epsilon_{t-1})$ for all t. Because g_t's are fixed vectors, the distribution of k_t is exactly the same whether or not we share the random variables across time points. With this reduction in number of random variables used, we now have

$$k_t = \begin{cases} k_{t-1} & \text{if } k_{t-1} \in \arg\max_{k=1}^{N} \sum_{\tau=1}^{t-1} (1+\epsilon_\tau) g_{\tau,k} \\ \in \arg\max_{k=1}^{N} \sum_{\tau=1}^{t-1} (1+\epsilon_\tau) g_{\tau,k} & \text{otherwise} \end{cases}. \quad (3)$$

We define $G_t = \sum_{\tau=1}^{t} g_\tau$ the cumulative payoff vector at time t. Define $\tilde{g}_t = \epsilon_t g_t$ and $\tilde{G}_t = \sum_{\tau=1}^{t} \tilde{g}_\tau$. We also define

$$g_{t,i-j} = g_{t,i} - g_{t,j}, \tilde{g}_{t,i-j} = \tilde{g}_{t,i} - \tilde{g}_{t,j}.$$

With these definitions, we have

$$\begin{aligned} \tilde{G}_{t,i-j} = \tilde{G}_{t,i} - \tilde{G}_{t,j} &= \sum_{\tau=1}^{t} \tilde{g}_{\tau,i} - \sum_{\tau=1}^{t} \tilde{g}_{\tau,j} \\ &= \sum_{\tau=1}^{t} (1+\epsilon_\tau)(g_{\tau,i} - g_{\tau,j}) = \sum_{\tau=1}^{t} (1+\epsilon_\tau) g_{\tau,i-j}. \end{aligned}$$

The following result upper bounds the regret in terms of downward zero-crossings of the process $\tilde{G}_{t,i-j}$, i.e., the times t when it switches from being non-negative at time $t-1$ to negative at time t.

Theorem 4.1. *We have the following upper bound on the expected regret:*

$$\mathbb{E}\left[\mathcal{R}_T\right] \leq 2N^2 \max_{1 \leq i,j \leq N} \sum_{t=1}^{T} |g_{t,i-j}| P\left(\tilde{G}_{t-1,i-j} \geq 0, \tilde{G}_{t,i-j} < 0\right).$$

We now focus on bounding the switching probabilities for a fixed pair i, j.

4.2 Step 2: Bounding Switching Probabilities Using Littlewood-Offord Theory

Our strategy is to do a "multi-scale" analysis and, within each scale, apply Littlewood-Offord theory to bound the switching probabilities. The need for a

multi-scale argument arises from the requirement in Littlewood-Offord theorem (see Theorem 4.2 below) for a lower bound on the step sizes of random walks. We partition the set of T time points $[T] := \{1, \ldots, T\}$ into $K+1$ disjoint sets at different scales, denoted as $\{A_k\}_{k=0}^{K}$ where

$$A_k = \begin{cases} \{t \in [T] : |g_{t,i-j}| \leq \frac{1}{\sqrt{T}}\} & k = 0 \\ \{t \in [T] : T^{-\frac{1}{2^k}} < |g_{t,i-j}| \leq T^{-\frac{1}{2^{k+1}}}\} & k = 1, \ldots, K-1 \\ \{t \in [T] : T^{-\frac{1}{2^K}} < |g_{t,i-j}| \leq 1\} & k = K \end{cases}.$$

The cardinality of a finite set A will be denoted by $|A|$. The number $K+1$ of different scales is determined by

$$K = \arg\min\{k \in \mathbb{N} \,:\, T^{-\frac{1}{2^k}} \geq 1/2\}.$$

It easily follows that,

$$\begin{aligned} &\sum_{t=1}^{T} |g_{t,i-j}| P\left(\tilde{G}_{t-1,i-j} \geq 0, \tilde{G}_{t,i-j} < 0\right) \\ =&\sum_{t=1}^{T} |g_{t,i-j}| P\left(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right) \\ =&\sum_{k=0}^{K} \sum_{t \in A_k} |g_{t,i-j}| P\left(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right). \end{aligned}$$

We now want to argue that the probabilities involved above are small. The crucial observation is that, if a switch occurs, then the random sum $\sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j}$ has to lie in a sufficiently small interval. Such "small ball" probabilities are exactly what the classic Littlewood-Offord theorem controls.

Theorem 4.2 (Littlewood-Offord Theorem of Erdös, Theorem 3 of Erdős [1945]). *Let $x_1, \ldots, x_n$ be n real numbers such that $|x_i| \geq 1$ for all i. For any given radius $\Delta > 0$, the small ball probability satisfies*

$$\sup_B P(\epsilon_1 x_1 + \cdots + \epsilon_n x_n \in B) \leq \frac{S(n)}{2^n}(\lfloor \Delta \rfloor + 1)$$

where $\epsilon_1, \ldots, \epsilon_n$ are i.i.d. Rademacher random variables, B ranges over all closed balls (intervals) of radius Δ, and $\lfloor x \rfloor$ refers to the integral part of x, $S(n)$ is the largest binomial coefficient belonging to n.

Using elementary calculations to upper bound $S(n)$ gives us the following corollary.

Corollary 4.2.1. *Under the same notation and conditions as Theorem 4.2, we have*

$$\sup_B P(\epsilon_1 x_1 + \cdots + \epsilon_n x_n \in B) \leq C_{LO}(\lfloor \Delta \rfloor + 1)\frac{1}{\sqrt{n}}$$

where $C_{LO} = \frac{2\sqrt{2}e}{\pi} < 3$.

The scale of payoffs for time periods in A_0 is so small that we do not need any Littlewood-Offord theory to control their contribution to the regret. Simply bounding the probabilities by 1 gives us the following.

Theorem 4.3. *The following upper bound holds for switching probabilities for time periods within A_0:*

$$\sum_{t\in A_0} |g_{t,i-j}| P\left(\sum_{\tau=1}^{t-1}\epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t}\epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right) \leq \sqrt{|A_0|}.$$

The real work lies in controlling the switching probabilities for payoffs at intermediate scales. The idea in the proof of the results below is to condition on the ϵ_t's outside A_k. Then the probability of interest is written as a small ball event in terms of the ϵ_t's in A_k. Applying Littlewood-Offord theorem then concludes the argument.

Theorem 4.4. *For any $k \in \{1, \ldots, K\}$, we have*

$$\sum_{t\in A_k} |g_{t,i-j}| P\left(\sum_{\tau=1}^{t-1}\epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t}\epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right) \leq 6C_{LO}\sqrt{|A_k|}.$$

We finally have all the ingredients in place to control the switching probabilities.

Corollary 4.4.1. *The following upper bound on the switching probabilities holds.*

$$\sum_{t=1}^{T} |g_{t,i-j}| P\left(\sum_{\tau=1}^{t-1}\epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t}\epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right) \leq 6C_{LO}\sqrt{T\log_2(4\log_2 T)}.$$

Proof. Using Theorem 4.3 and Theorem 4.4, we have

$$\begin{aligned}
&\sum_{t=1}^{T} |g_{t,i-j}| P\left(\sum_{\tau=1}^{t-1}\epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t}\epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right) \\
=&\sum_{k=0}^{K}\sum_{t\in A_k} |g_{t,i-j}| P\left(\sum_{\tau=1}^{t-1}\epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t}\epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right) \\
\leq&\sum_{k=0}^{K} 6C_{LO}\sqrt{|A_k|}.
\end{aligned}$$

Since $\sum_k \sqrt{|A_k|} \leq \sqrt{K+1}\cdot\sqrt{\sum_k |A_k|}$ and $\sum_{k=0}^{K}|A_k| = T$, we have

$$\sum_{k=0}^{K} 6C_{LO}\sqrt{|A_k|} \leq 6C_{LO}\sqrt{(K+1)T}.$$

By definition of K, we have that $T^{-\frac{1}{2^{K-1}}} < \frac{1}{2}$, $K \leq \log_2(\log_2(T)) + 1$ which finishes the proof. □

Thus, $\forall i, j \in \{1, \ldots, N, i \neq j\}$, we have

$$\sum_{t=1}^{T} |g_{t,i-j}| P\left(\tilde{G}_{t-1,i-j} \geq 0, \tilde{G}_{t,i-j} < 0\right) \leq 6C_{LO}\sqrt{T \log_2(4 \log_2 T)},$$

which, when plugged into Theorem 4.1, immediately yields the following corollary.

Corollary 4.4.2. *Against an oblivious opponent, both versions — the single stream version (3) and the fresh-randomization-at-each-round version (1) — of sampled fictitious play enjoy the following bound on expected regret.*

$$\mathbb{E}[\mathcal{R}_T] \leq 12C_{LO}N^2\sqrt{T \log_2(4 \log_2 T)}.$$

4.3 Step 3: From Oblivious to Adaptive Opponents

Now we consider adaptive opponents. In this setting, we can no longer assume that player i plays against a fix sequence of payoff vectors g_t. Note that $g_{t,k}$ is just shorthand for $u_i(k, s_{t,-i})$ and opponents can react to player i's moves $k_1, \ldots, k_{t-1}$ in selecting their strategy tuple $s_{t,-i}$, possibly making use of their own private stream of i.i.d. random variables $U_1, U_2, \ldots$. Thus g_t is a function $g_t(k_1, \ldots, k_{t-1}, U_t)$. Faced with general adaptive opponents, the single stream version (3) can incur terrible regret as stated below.

Theorem 4.5. *The single stream version of the sampled fictitious play procedure (3) can incur linear expected regret against adaptive opponents.*

However, for the fresh randomness at each round procedure (1), we can apply Lemma 4.1 of Cesa-Bianchi and Lugosi [2006] along with Corollary 4.4.2 to derive our next result that holds for adaptive opponents too. There are two conditions that we must verify before we apply that lemma. First, the learning procedure should use independent randomization at different time points. Second, the probability distribution of $s_{t,i}$ over the N available strategies should be fully determined by $s_{1,-i}, \ldots, s_{t-1,-i}$ and should not depend explicitly on player i own previous moves $s_{1,i}, \ldots, s_{t-1,i}$. Both of these conditions are easily seen to hold for sampled fictitious play as defined in (1) and (2).

Theorem 4.6. *For any $\delta_t > 0$, with probability at least $1 - \delta_t$, the actual regret $\mathcal{R}_t$ of sampled fictitious play as defined in (1) and (2) satisfies, for any adaptive opponent,*

$$\mathcal{R}_t \leq 12C_{LO}N^2\sqrt{t \log_2(4 \log_2 t)} + \sqrt{\frac{t}{2} \log \frac{1}{\delta_t}}.$$

Now fix an arbitrary $\epsilon > 0$. Consider the events $E_t = \{\mathcal{R}_t \geq \epsilon t\}$. Setting $\delta_t = 1/t^2$ in the above theorem, gives us $P(E_t) \leq \delta_t$ for large enough t. Since $\sum_t \delta_t < \infty$, we have $\sum_t P(E_t) < \infty$. Therefore, using Borel-Cantelli lemma, the event "infinitely many E_t's occur" has probability 0. That is, with probability 1, we have $\limsup_{t\to\infty} \frac{\mathcal{R}_t}{t} \leq \epsilon$. Since $\epsilon > 0$ was arbitrary, this proves Theorem 3.1.

5 Conclusion

We proved that a natural variant of fictitious play is Hannan consistent. In the variant we considered, the player plays best response to moves of her opponents at sampled time points in the history so far. We considered one particular sampling scheme, namely Bernoulli sampling. It will be interesting to consider other sampling strategies including sampling with replacement. It will also be interesting to consider notions of regret, such as tracking regret [Cesa-Bianchi and Lugosi, 2006, Section 5.2], that are more suitable for non-stationary environments by biasing the sampling to give more importance to recent time points.

Acknowledgements

We thank Jacob Abernethy, Gergely Neu and Manfred Warmuth for helpful discussions. We acknowledge the support of NSF via CAREER grant IIS-1452099.

References

Avrim Blum and Yishay Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8(Jun):1307–1324, 2007.

Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

Luc Devroye, Gábor Lugosi, and Gergely Neu. Prediction by random-walk perturbation. In *COLT*, pages 460–473, 2013.

Paul Erdős. On a lemma of Littlewood and Offord. *Bulletin of the American Mathematical Society*, 51:898–902, 1945.

Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*. MIT Press, 1998.

Dennis Gilliland and Inha Jung. Play against the random past for matching binary bits. *Journal of Statistical Theory and Application*, 5(3):282–291, 2006.

James Hannan. Approximation to Bayes risk in repeated play. *Contributions to the theory of games*, 3(39):97–139, 1957.

Sergiu Hart and Andreu Mas-Colell. *Simple Adaptive Strategies: From Regret Matching to Uncoupled Dynamics*, volume 4 of *World Scientific Series in Economic Theory*. World Scientific Publishing, 2013.

Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.

Yuri M. Kaniovski and H. Peyton Young. Learning dynamics in games with stochastic perturbations. *Games and Economic Behavior*, 11(2):330–363, 1995.

Theodore J. Lambert III, Marina A. Epelman, and Robert L. Smith. A fictitious play approach to large-scale optimization. *Operation Research*, 53(3):477–489, 2005.

Gilles Stoltz and Gábor Lugosi. Internal regret in on-line portfolio selection. *Machine Learning*, 59(1-2):125–159, 2005.

Tim van Erven, Wojciech Kotlowski, and Manfred K. Warmuth. Follow the leader with dropout perturbations. In *Proceedings of Conference on Learning Theory (COLT)*, pages 949–974, 2014.

Appendix A Proofs

Proof of Theorem 4.1. We will prove a result for Bernoulli sampling with general probabilities, i.e., when $P(\epsilon_t = +1) = \alpha$ where α is not necessarily $1/2$. We will show that

$$\mathbb{E}\left[\mathcal{R}_T\right] \leq \frac{N^2}{\alpha} \max_{1 \leq i,j \leq N} \sum_{t=1}^{T} |g_{t,i-j}| P\left(\tilde{G}_{t-1,i-j} \geq 0, \tilde{G}_{t,i-j} < 0\right)$$

from which the theorem follows as a special case when $\alpha = 1/2$.

Obviously we have $\mathbb{E}(\tilde{g}_{t,i}) = 2\alpha g_{t,i}$ because of the fact that $\mathbb{E}(\epsilon_t) = 2\alpha - 1$. Furthermore, $\mathbb{E}[\tilde{g}_{t,k_t} | \epsilon_1, \ldots, \epsilon_{t-1}] = 2\alpha g_{t,k_t}$, which implies that $\mathbb{E}[\tilde{g}_{t,k_t}] = \mathbb{E}\left[\mathbb{E}[\tilde{g}_{t,k_t} | \epsilon_1, \ldots, \epsilon_{t-1}]\right] = 2\alpha \mathbb{E}[g_{t,k_t}]$. We now have,

$$\begin{aligned}
\mathbb{E}\left[\mathcal{R}_T\right] &= \max_{k=1}^{N} \sum_{t=1}^{T} g_{t,k} - \mathbb{E}\left[\sum_{t=1}^{T} g_{t,k_t}\right] \\
&= \frac{1}{2\alpha} \max_{k=1}^{N} \mathbb{E}\left[\sum_{t=1}^{T} \tilde{g}_{t,k}\right] - \frac{1}{2\alpha} \mathbb{E}\left[\sum_{t=1}^{T} \tilde{g}_{t,k_t}\right] \\
&\leq \frac{1}{2\alpha} \mathbb{E}\left[\max_{k=1}^{N} \sum_{t=1}^{T} \tilde{g}_{t,k} - \sum_{t=1}^{T} \tilde{g}_{t,k_t}\right].
\end{aligned}$$

Using Lemma A.1, we can further upper bound the last expression as follows,

$$\begin{aligned}
&\leq \frac{1}{2\alpha} \mathbb{E}\left[\sum_{t=1}^{T} \tilde{g}_{t,k_{t+1}} - \sum_{t=1}^{T} \tilde{g}_{t,k_t}\right] \\
&= \frac{1}{2\alpha} \sum_{t=1}^{T} \mathbb{E}\left[(1+\epsilon_t)(g_{t,k_{t+1}} - g_{t,k_t})\right] \\
&\leq \frac{1}{2\alpha} \sum_{t=1}^{T} \mathbb{E}\left[(1+\epsilon_t)|g_{t,k_{t+1}} - g_{t,k_t}|\right] \\
&\leq \frac{1}{\alpha} \sum_{t=1}^{T} \mathbb{E}\left[|g_{t,k_t} - g_{t,k_{t+1}}|\right] \\
&= \frac{1}{\alpha} \sum_{t=1}^{T} \sum_{1 \leq i,j \leq N} \mathbb{E}\left[|g_{t,i} - g_{t,j}| 1_{(k_t = i, k_{t+1} = j)}\right] \\
&= \frac{1}{\alpha} \sum_{1 \leq i,j \leq N} \sum_{t=1}^{T} \mathbb{E}\left[|g_{t,i} - g_{t,j}| 1_{(k_t = i, k_{t+1} = j)}\right] \\
&\leq \frac{N^2}{\alpha} \max_{1 \leq i,j \leq N} \sum_{t=1}^{T} |g_{t,i} - g_{t,j}| P(k_t = i, k_{t+1} = j) \\
&\leq \frac{N^2}{\alpha} \max_{1 \leq i,j \leq N} \sum_{t=1}^{T} |g_{t,i} - g_{t,j}| P\left(\tilde{G}_{t-1,i} \geq \tilde{G}_{t-1,j}, \tilde{G}_{t,i} < \tilde{G}_{t,j}\right) \\
&\leq \frac{N^2}{\alpha} \max_{1 \leq i,j \leq N} \sum_{t=1}^{T} |g_{t,i-j}| P\left(\tilde{G}_{t-1,i-j} \geq 0, \tilde{G}_{t,i-j} < 0\right).
\end{aligned}$$

□

Lemma A.1.

$$\sum_{t=1}^{T} \tilde{g}_{t,k_{t+1}} \geq \sum_{t=1}^{T} \tilde{g}_{t,k_{T+1}} = \max_{k=1}^{N} \sum_{t=1}^{T} \tilde{g}_{t,k}$$

Proof. The proof goes by induction. The statement is obvious for T = 1. Assume now that

$$\sum_{t=1}^{n-1} \tilde{g}_{t,k_{t+1}} \geq \sum_{t=1}^{n-1} \tilde{g}_{t,k_n}.$$

Since by definition $\sum_{t=1}^{n-1} \tilde{g}_{t,k_n} \geq \sum_{t=1}^{n-1} \tilde{g}_{t,k_{n+1}}$, the inductive assumption implies

$$\sum_{t=1}^{n-1} \tilde{g}_{t,k_{t+1}} \geq \sum_{t=1}^{n-1} \tilde{g}_{t,k_{n+1}}.$$

Add $\tilde{g}_{n,k_{n+1}}$ to both sides to obtain the result. □

Proof of Corollary 4.2.1. Note that when $\alpha = \frac{1}{2}$ Lemma A.2 provides a bound on $\frac{S(n)}{2^n}$. Plug in $\alpha = \frac{1}{2}$ to Lemma A.2 and combine with Theorem 4.2, we know that if $n > 4$, $C_{LO} = \frac{\sqrt{2}e}{\pi}$ will suffice. If $n \leq 4$, $\frac{2\sqrt{2}e}{\pi} \times n^{-\frac{1}{2}} > 1$ and Lemma A.2 holds. □

Lemma A.2. *Suppose $X_1, \ldots, X_t$ are i.i.d. Bernoulli random variables that take value of 1 with probability α and 0 with probability $1-\alpha$. If $t > \max(\frac{2}{1-\alpha}, \frac{2}{\alpha}) \geq \max(\frac{2\alpha}{1-\alpha}, \frac{2}{\alpha})$, then for all a,*

$$P\left(\sum_{i=1}^{t} X_i = a\right) \leq \frac{e}{2\pi} \times \sqrt{\frac{2}{\alpha(1-\alpha)}} \times t^{-\frac{1}{2}}.$$

Proof. It has been shown that the maximum probability of Bernoulli distribution $P(X = a)$ is achieved when $a = \hat{a} = \lfloor (t+1)\alpha \rfloor$ where $\lfloor x \rfloor$ denotes the integral part of x. Clearly $\hat{a} \in [t\alpha - 1, (t+1)\alpha]$. Thus,

$$\begin{aligned}
\sqrt{\hat{a}(t-\hat{a})} &\geq \min\left(\sqrt{(t\alpha-1)(t-t\alpha+1)}, \sqrt{(t+1)\alpha(t-t\alpha-\alpha)}\right) \\
&= t \times \min\left(\sqrt{(\alpha - \frac{1}{t})(1-\alpha+\frac{1}{t})}, \sqrt{(1+\frac{1}{t})\alpha(1-\alpha-\frac{\alpha}{t})}\right) \\
&\geq t \times \min\left(\sqrt{(\alpha - \frac{\alpha}{2})(1-\alpha)}, \sqrt{\alpha(1-\alpha-\frac{1-\alpha}{2})}\right) \\
&= \sqrt{\frac{\alpha(1-\alpha)}{2}} t.
\end{aligned}$$

With this preliminary inequality, we are ready to prove the lemma.

$$\begin{aligned}
P(\sum_{i=1}^{t} X_i = a) &\leq P(\sum_{i=1}^{t} X_i = \hat{a}) \\
&= \binom{t}{\hat{a}} \times \alpha^{\hat{a}} (1-\alpha)^{t-\hat{a}}
\end{aligned}$$

$$
\begin{aligned}
&= \frac{t!}{(\hat{a})!(t-\hat{a})!} \times \alpha^{\hat{a}}(1-\alpha)^{t-\hat{a}} \\
&\leq \frac{t^{t+\frac{1}{2}}e^{1-t}}{\left(\sqrt{2\pi}\hat{a}^{\hat{a}+\frac{1}{2}}e^{-\hat{a}}\right)\left(\sqrt{2\pi}(t-\hat{a})^{t-\hat{a}+\frac{1}{2}}e^{-(t-\hat{a})}\right)} \times \alpha^{\hat{a}}(1-\alpha)^{t-\hat{a}} \\
&= \frac{e}{2\pi} \times \frac{1}{\sqrt{\hat{a}(t-\hat{a})}} \times \frac{t^{t+\frac{1}{2}}}{\hat{a}^{\hat{a}}(t-\hat{a})^{t-\hat{a}}} \times \alpha^{\hat{a}}(1-\alpha)^{t-\hat{a}} \\
&\leq \frac{e}{2\pi} \times \sqrt{\frac{2}{\alpha(1-\alpha)}} \times t^{t-\frac{1}{2}} \times \frac{\alpha^{\hat{a}}(1-\alpha)^{t-\hat{a}}}{\hat{a}^{\hat{a}}(t-\hat{a})^{t-\hat{a}}}.
\end{aligned}
$$

Let $f(x) = \frac{\alpha^x(1-\alpha)^{t-x}}{x^x(t-x)^{t-x}}$, $f'(x) = \left(log(\frac{\alpha}{1-\alpha}) - log(\frac{x}{t-x})\right) \times f(x)$. Obviously $f'(x)$ is 0 when $x = \alpha t$, positive when $x < \alpha t$, and negative when $x > \alpha t$. Thus,

$$
f(x) \leq \frac{\alpha^{\alpha t}(1-\alpha)^{t-\alpha t}}{(\alpha t)^{\alpha t}(t-\alpha t)^{t-\alpha t}} = t^{-t}.
$$

Hence,

$$
\begin{aligned}
P(\sum_{i=1}^{t} X_i = a) &\leq \frac{e}{2\pi} \times \sqrt{\frac{2}{\alpha(1-\alpha)}} \times t^{t-\frac{1}{2}} \times f(\hat{a}) \\
&\leq \frac{e}{2\pi} \times \sqrt{\frac{2}{\alpha(1-\alpha)}} \times t^{-\frac{1}{2}}.
\end{aligned}
$$

□

Proof of Theorem 4.3. We write $|A|$ to denote the cardinality of a finite set A.

$$
\begin{aligned}
&\sum_{t\in A_0} |g_{t,i-j}| P(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}) \\
\leq &\sum_{t\in A_0} \frac{1}{\sqrt{T}} \times 1 = \frac{|A_0|}{\sqrt{T}} \leq \sqrt{|A_0|}.
\end{aligned}
$$

□

Proof of Theorem 4.4. We write ϵ with a subset of $[T]$ as subscript to denote ϵ_t's at times that are within the subset. For example, $\epsilon_{[T]} = \{\epsilon_1, \ldots, \epsilon_T\}$. We also write ϵ_{-A} to denote the set of ϵ_t's that are within the complement of A with respect to $[T]$.

Case I: $k \in \{1, \ldots, K-1\}$

$$
\begin{aligned}
&\sum_{t\in A_k} |g_{t,i-j}| P\left(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right) \\
=&\sum_{t\in A_k} |g_{t,i-j}| \mathbb{E}_{\epsilon_{[T]}}\left[\mathbb{1}_{(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j})}\right] \\
=&\sum_{t\in A_k} |g_{t,i-j}| \mathbb{E}_{\epsilon_{-A_k}}\left[\mathbb{E}_{\epsilon_{A_k}}\left[\mathbb{1}_{(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j})} | \epsilon_{-A_k}\right]\right] \\
=&\mathbb{E}_{\epsilon_{-A_k}}\left[\sum_{t\in A_k} |g_{t,i-j}| \mathbb{E}_{\epsilon_{A_k}}\left[\mathbb{1}_{(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j})} | \epsilon_{-A_k}\right]\right]
\end{aligned}
$$

$$\leq \sup_{\epsilon_{-A_k}} \sum_{t\in A_k} |g_{t,i-j}| \mathbb{E}_{\epsilon_{A_k}} [\mathbb{1}_{(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j})} | \epsilon_{-A_k}].$$

Let $A_k = \{t_{k,1}, \ldots, t_{k,|A_k|}\}$ with elements listed in increasing order of time index. Then, we have

$$\begin{aligned}
&\sum_{t\in A_k} |g_{t,i-j}| \mathbb{E}_{\epsilon_{A_k}} [\mathbb{1}_{(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j})} | \epsilon_{-A_k}] \\
=&\sum_{n=1}^{|A_k|} |g_{t_{k,n},i-j}| \mathbb{E}_{\epsilon_{A_k}} [\mathbb{1}_{(\sum_{s=1}^{n-1} \epsilon_{t_{k,s}} g_{t_{k,s},i-j} \geq -\sum_{s=1}^{n-1} g_{t_{k,s},i-j} + D_n, \sum_{s=1}^{n} \epsilon_{t_{k,s}} g_{s,i-j} < -\sum_{s=1}^{n} g_{t_{k,s},i-j} + D_n)} | \epsilon_{-A_k}] \\
=&\sum_{n=1}^{|A_k|} |g_{t_{k,n},i-j}| P \left(\sum_{s=1}^{n-1} \epsilon_{t_{k,s}} g_{t_{k,s},i-j} \geq -\sum_{s=1}^{n-1} g_{t_{k,s},i-j} + D_n, \right. \\
&\qquad \left. \sum_{s=1}^{n} \epsilon_{t_{k,s}} g_{t_{k,s},i-j} < -\sum_{s=1}^{n} g_{t_{k,s},i-j} + D_n \middle| \epsilon_{-A_k} \right)
\end{aligned}$$

where

$$D_n = D_n(\epsilon_{-A_k}) = -\sum_{\tau=1, \tau\in -A_k}^{t_{k,n}-1} (1+\epsilon_\tau) g_{\tau,i-j}.$$

By definition of the set A_k, we have $|g_{t_{k,s},i-j}| \geq T^{-\frac{1}{2^k}}$, so $T^{\frac{1}{2^k}} |g_{t_{k,s},i-j}| \geq 1$. Let $M_k = T^{\frac{1}{2^k}}$. Then, we have

$$\begin{aligned}
&\sum_{n=1}^{|A_k|} |g_{t_{k,n},i-j}| P \left(\sum_{s=1}^{n-1} \epsilon_{t_{k,s}} g_{t_{k,s},i-j} \geq -\sum_{s=1}^{n-1} g_{t_{k,s},i-j} + D_n, \right. \\
&\qquad \left. \sum_{s=1}^{n} \epsilon_{t_{k,s}} g_{t_{k,s},i-j} < -\sum_{s=1}^{n} g_{t_{k,s},i-j} + D_n \middle| \epsilon_{-A_k} \right) \\
=&\sum_{n=1}^{|A_k|} |g_{t_{k,n},i-j}| P \left(\sum_{s=1}^{n-1} \epsilon_{t_{k,s}} g_{t_{k,s},i-j} M_k \geq -\sum_{s=1}^{n-1} g_{t_{k,s},i-j} M_k + D_n M_k , \right. \\
&\qquad \left. \sum_{s=1}^{n} \epsilon_{t_{k,s}} g_{t_{k,s},i-j} M_k < -\sum_{s=1}^{n} g_{t_{k,s},i-j} M_k + D_n M_k \middle| \epsilon_{-A_k} \right) \\
\leq&\sum_{n=1}^{|A_k|} |g_{t_{k,n},i-j}| P \left(\sum_{s=1}^{n} \epsilon_{t_{k,s}} g_{t_{k,s},i-j} M_k \in B_{k,n} \middle| \epsilon_{-A_k} \right)
\end{aligned}$$

where

$$B_{k,n} = \left[-\sum_{s=1}^{n} g_{t_{k,s},i-j} M_k + D_n M_k - 2|g_{t_{k,n},i-j}| M_k, -\sum_{s=1}^{n} g_{t_{k,s},i-j} M_k + D_n M_k \right]$$

is a one-dimensional closed ball with radius $\Delta = |g_{t_{k,n},i-j}| M_k$. Note that this ball is fixed given ϵ_{-A_k}. Since $|g_{t_{k,s},i-j}| M_k \geq 1$, we can apply Corollary 4.2.1 to get

$$P \left(\sum_{s=1}^{n} \epsilon_{t_{k,s}} g_{t_{k,s},i-j} M_k \in B_{k,n} \middle| \epsilon_{-A_k} \right) \leq \frac{C_{LO}(\Delta+1)}{\sqrt{n}} = \frac{C_{LO}(|g_{t_{k,n},i-j}| M_k + 1)}{\sqrt{n}}.$$

Now we continue the derivation,

$$\begin{aligned}
&\sum_{n=1}^{|A_k|} |g_{t_{k,n},i-j}| P\left(\sum_{s=1}^{n} \epsilon_{t_{k,s}} g_{t_{k,s},i-j} M_k \in B_{k,n} \middle| \epsilon_{-A_k}\right) \\
\leq &\sum_{n=1}^{|A_k|} |g_{t_{k,n},i-j}| \frac{C_{LO}(|g_{t_{k,n},i-j}|M_k + 1)}{\sqrt{n}} \\
\leq &C_{LO}\left(\sum_{n=1}^{|A_k|} \frac{|g_{t_{k,n},i-j}|^2 M_k}{\sqrt{n}} + \sum_{n=1}^{|A_k|} \frac{1}{\sqrt{n}}\right).
\end{aligned}$$

Since we have $|g_{t_{k,n},i-j}| < T^{-\frac{1}{2^{k+1}}}$, $|g_{t_{k,n},i-j}|^2 T^{\frac{1}{2^k}} = |g_{t_{k,n},i-j}|^2 M_k < 1$. Thus we have the bound,

$$\begin{aligned}
&C_{LO}\left(\sum_{n=1}^{|A_k|} \frac{|g_{t_{k,n},i-j}|^2 M_k}{\sqrt{n}} + \sum_{n=1}^{|A_k|} \frac{1}{\sqrt{n}}\right) \\
\leq &2C_{LO} \sum_{n=1}^{|A_k|} \frac{1}{\sqrt{n}} \leq 4C_{LO}\sqrt{|A_k|}.
\end{aligned}$$

Case II: $k = K$. Similar to the previous case, we have

$$\begin{aligned}
&\sum_{t\in A_K} |g_{t,i-j}| P\left(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right) \\
\leq &\sup_{\epsilon_{-A_K}} \sum_{t\in A_K} |g_{t,i-j}| \mathbb{E}_{\epsilon_{A_k}}[\mathbb{1}_{(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j})} | \epsilon_{-A_K}]
\end{aligned}$$

and writing the elements of A_K in increasing order as $\{t_{K,1}, \ldots, t_{K,|A_K|}\}$, we get

$$\begin{aligned}
&\sum_{t\in A_K} |g_{t,i-j}| \mathbb{E}_{\epsilon_{A_K}}[\mathbb{1}_{(\sum_{\tau=1}^{t-1} \epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t} \epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j})} | \epsilon_{-A_K}] \\
\leq &\sum_{n=1}^{|A_K|} |g_{t_{K,n},i-j}| P\left(\sum_{s=1}^{n} \epsilon_{t_{K,s}} g_{t_{K,s},i-j} M_K \in B_{K,n}\right)
\end{aligned}$$

where

$$D_n = D_n(\epsilon_{-A_K}) = -\sum_{\tau=1, \tau\in\{-A_K\}}^{t_{K,n}-1} (1 + \epsilon_\tau) g_{\tau,i-j},$$

$M_K = T^{\frac{1}{2^K}} \leq 2$, and

$$B_{K,n} = \left[-\sum_{s=1}^{n} g_{t_{K,s},i-j} M_K + D_n M_K - 2|g_{t_{K,n},i-j}| M_K, -\sum_{s=1}^{n} g_{t_{K,s},i-j} M_K + D_n M_K\right]$$

is a one-dimensional closed ball with radius $\Delta = |g_{t_{K,n},i-j}| M_K$. Note that this ball is fixed given ϵ_{-A_K} and hence, we can apply Corollary 4.2.1 to get

$$\sum_{n=1}^{|A_K|} |g_{t_{K,n},i-j}| P\left(\sum_{s=1}^{n} \epsilon_{t_{K,s}} g_{t_{K,s},i-j} M_K \in B_{K,n}\right)$$

$$\leq \sum_{n=1}^{|A_K|} |g_{t_{K,n},i-j}| \frac{C_{LO}(|g_{t_{K,n},i-j}|M_K + 1)}{\sqrt{n}}$$

$$\leq C_{LO}(M_K + 1) \sum_{n=1}^{|A_K|} \frac{1}{\sqrt{n}} \leq 6C_{LO}\sqrt{|A_K|}.$$

Combining the two cases proves the theorem. □

Proof of Theorem 4.5. Consider a game with two strategies, i.e., $N = 2$. We refer to player i as the "player" and the other players collectively as the "environment". On odd rounds, the environment plays payoff vector $(0, 0)$. This ensures that after odd rounds, the environment will know exactly which strategy the player will choose, because no matter whether the Rademacher random variable is -1 or $+1$, the next strategy played will be the same as the strategy the player just played. On even rounds, the environment plays $(0, 1)$ if the player chose the first strategy in the previous round, and $(1, 0)$ if the player chose the second strategy in the previous round. The payoff acquired by the player by following sampled fictitious play procedure will be 0 because the player gains 0 on all rounds. However, as evident from the environment's procedure, the total payoff for two strategies is $T/2$ and thus the best strategy has a payoff no less than $T/4$ because of the pigeonhole principle. Hence, the expected regret for the player is at least $T/4$. □

Appendix B Asymmetric Probabilities

In this section we would like to prove that for binary payoff and $\alpha \in (0,1)$ insead of just $1/2$, the expected regret is $O(\sqrt{T})$ where the constant hidden in $O(\cdot)$ notation blows up in either of the two extreme case: $\alpha \to 0$ and $\alpha \to 1$. Note that we are still considering the single stream version (3) of the learning procedure.

Theorem B.1. *For $\alpha \in (0,1)$ and $g_t \in \{0,1\}^N$, assuming that $T > \max(\frac{2}{1-\alpha}, \frac{2}{\alpha})$, the expected regret satisfies*

$$\mathbb{E}\left[\mathcal{R}_T\right] \leq \frac{4N^2 Q_\alpha}{\alpha}\sqrt{T}$$

where $Q_\alpha = \frac{e}{2\pi} \times \sqrt{\frac{2}{\alpha(1-\alpha)}}$.

Proof. We begin with the inequality obtained in the proof of Theorem 4.1:

$$\mathbb{E}\left[\mathcal{R}_T\right] \leq \frac{N^2}{\alpha} \max_{1\leq i,j\leq N} \sum_{t=1}^{T} |g_{t,i-j}| P\left(\tilde{G}_{t-1,i-j} \geq 0, \tilde{G}_{t,i-j} < 0\right). \tag{4}$$

As before, we fix i and j, and will bound the expression

$$\sum_{t=1}^{T} |g_{t,i-j}| P(\sum_{\tau=1}^{t-1}(1+\epsilon_\tau)g_{\tau,i-j} \geq 0, \sum_{\tau=1}^{t}(1+\epsilon_\tau)g_{\tau,i-j} < 0).$$

The rest of the proof is similar to the proof of Theorem 4.4. Define the classes $A_1 = \{t : g_{t,i-j} = 1, t = 1,\ldots,T\}$ and $A_{-1} = \{t : g_{t,i-j} = -1, t = 1,\ldots,T\}$. We have,

$$\begin{aligned}
&\sum_{t=1}^{T} |g_{t,i-j}| P\left(\sum_{\tau=1}^{t-1}(1+\epsilon_\tau)g_{\tau,i-j} \geq 0, \sum_{\tau=1}^{t}(1+\epsilon_\tau)g_{\tau,i-j} < 0\right) \\
\leq &\sum_{t=1}^{T} P\left(\sum_{\tau=1}^{t-1}\epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t}\epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right) \\
= &\sum_{k\in\{\pm 1\}} \sum_{t\in A_k} P\left(\sum_{\tau=1}^{t-1}\epsilon_\tau g_{\tau,i-j} \leq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t}\epsilon_s g_{\tau,i-j} > -\sum_{\tau=1}^{t} g_{\tau,i-j}\right).
\end{aligned}$$

For any $k \in \{\pm 1\}$,

$$\begin{aligned}
&\sum_{t\in A_k} P\left(\sum_{\tau=1}^{t-1}\epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t}\epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}\right) \\
\leq &\sup_{\epsilon_{-A_k}} \sum_{t\in A_k} \mathbb{E}_{\epsilon_{A_k}}[\mathbb{1}_{\sum_{\tau=1}^{t-1}\epsilon_\tau g_{\tau,i-j} \geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t}\epsilon_\tau g_{\tau,i-j} < -\sum_{\tau=1}^{t} g_{\tau,i-j}} | \epsilon_{-A_k}].
\end{aligned}$$

Let $A_k = \{t_{k,1},\ldots,t_{k,|A_k|}\}$ with elements listed in increasing order of time index. Also define, for $n \in \{1,\ldots,|A_k|\}$,

$$D_n = D_n(\epsilon_{-A_k}) = -\sum_{\tau=1,\tau\in\{-A_k\}}^{t_{k,n}-1} \epsilon_\tau g_{\tau,i-j} - \sum_{\tau=1,\tau\in\{-A_k\}}^{t_{k,n}-1} g_{\tau,i-j}.$$

We then proceed as follows.

$$
\begin{aligned}
&\sum_{t\in A_k} \mathbb{E}_{\epsilon_{A_k}}\left[\mathbb{1}_{\sum_{\tau=1}^{t-1}\epsilon_\tau g_{\tau,i-j}\geq -\sum_{\tau=1}^{t-1} g_{\tau,i-j}, \sum_{\tau=1}^{t}\epsilon_\tau g_{\tau,i-j}< -\sum_{\tau=1}^{t} g_{\tau,i-j}} \middle| \epsilon_{-A_k}\right] \\
=&\sum_{n=1}^{|A_k|} \mathbb{E}_{\epsilon_{A_k}}\left[\mathbb{1}_{(\sum_{s=1}^{n-1}\epsilon_{t_{k,s}} g_{t_{k,s},i-j}\geq -\sum_{s=1}^{n-1} g_{t_{k,s},i-j}+D_n, \sum_{s=1}^{n}\epsilon_{t_{k,s}} g_{s,i-j}< -\sum_{s=1}^{n} g_{t_{k,s},i-j}+D_n)} \middle| \epsilon_{-A_k}\right] \\
=&\sum_{n=1}^{|A_k|} P\left(\sum_{s=1}^{n-1}\epsilon_{t_{k,s}} g_{t_{k,s},i-j}\geq -\sum_{s=1}^{n-1} g_{t_{k,s},i-j}+D_n,\right. \\
&\qquad \left.\sum_{s=1}^{n}\epsilon_{t_{k,s}} g_{t_{k,s},i-j}< -\sum_{s=1}^{n} g_{t_{k,s},i-j}+D_n \middle| \epsilon_{-A_k}\right) \\
\leq&\sum_{n=1}^{|A_k|} P\left(\sum_{s=1}^{n}\epsilon_{t_{k,s}} g_{t_{k,s},i-j}= -\sum_{s=1}^{n-1} g_{t_{k,s},i-j}+D_n-2 \vee\right. \\
&\qquad \left.\sum_{s=1}^{n}\epsilon_{t_{k,s}} g_{t_{k,s},i-j}= -\sum_{s=1}^{n-1} g_{t_{k,s},i-j}+D_n-1 \middle| \epsilon_{-A_k}\right) \\
\leq&\sum_{n=1}^{|A_k|} \left(P\left(\sum_{s=1}^{n}\epsilon_{t_{k,s}} g_{t_{k,s},i-j}= -\sum_{s=1}^{n-1} g_{t_{k,s},i-j}+D_n-2 \middle| \epsilon_{-A_k}\right)+\right. \\
&\qquad \left. P\left(\sum_{s=1}^{n}\epsilon_{t_{k,s}} g_{t_{k,s},i-j}= -\sum_{s=1}^{n-1} g_{t_{k,s},i-j}+D_n-1 \middle| \epsilon_{-A_k}\right)\right) \\
\leq&2\sum_{n=1}^{|A_k|} \frac{Q_\alpha}{\sqrt{n}} \leq 4Q_\alpha\sqrt{|A_k|} \leq 4Q_\alpha\sqrt{T}
\end{aligned}
$$

where $Q_\alpha = \frac{e}{2\pi}\times\sqrt{\frac{2}{\alpha(1-\alpha)}}$ from Lemma A.2. Putting this bound together with (4) concludes the proof. □